DESQ: Frequent Sequence Mining with Subsequence Constraints

Kaustubh Beedkar
University of Mannheim, Germany
Email: kbeedkar@uni-mannheim.de

Rainer Gemulla
University of Mannheim, Germany
Email: rgemulla@uni-mannheim.de

arXiv:1609.08431v1 [cs.DB] 27 Sep 2016

***Abstract*—Frequent sequence mining methods often make use of constraints to control which subsequences should be mined. A variety of such subsequence constraints has been studied in the literature, including length, gap, span, regular-expression, and hierarchy constraints. In this paper, we show that many subsequence constraints—including and beyond those considered in the literature—can be unified in a single framework. A unified treatment allows researchers to study jointly many types of subsequence constraints (instead of each one individually) and helps to improve usability of pattern mining systems for practitioners. In more detail, we propose a set of simple and intuitive "pattern expressions" to describe subsequence constraints and explore algorithms for efficiently mining frequent subsequences under such general constraints. Our algorithms translate pattern expressions to compressed finite state transducers, which we use as computational model, and simulate these transducers in a way suitable for frequent sequence mining. Our experimental study on real-world datasets indicates that our algorithms—although more general—are competitive to existing state-of-the-art algorithms.**

I. Introduction

Frequent sequence mining (FSM) is a fundamental task in data mining. Frequent sequences are useful for a wide range of applications, including market basket analysis [1], web usage mining and session analysis [2], natural language processing [3], information extraction [4], [5], or computational biology [6]. In web usage mining, for example, frequent sequences describe common behavior across users (e.g., the order in which users visit web pages). As another example, frequent textual patterns such as "*PERSON* is married to *PERSON*" are indicative of typed relations between entities and useful for natural-language processing and information extraction tasks [4], [5].

In FSM, we model the available data as a collection of sequences composed of items such as words (text processing), products (market-basket analysis), or actions and events (session analysis). Often items are arranged in an application-specific hierarchy; e.g., *is→be→VERB* (for words), *Canon 5D→DSLR camera→electronics* (for products), or *Rakesh Agrawal→scientist→PERSON* (for entities). The goal of FSM is to discover subsequences or generalized subsequences that occur in sufficiently many input sequences. Since the total number of such subsequences can potentially be very large and not all frequent subsequences may be of interest to a particular application, most FSM methods make use of subsequence constraints to control the set of subsequences to be mined.

A large variety of subsequence constraints has been studied in prior work [1], [7]–[13]. Commonly proposed constraints include *gap or span constraints*, where items in the subsequences need to appear "close" in the input sequence, and *length constraints*, where the number of items in the subsequences is bounded. In n-gram mining [14], for example, the goal is to mine frequent consecutive subsequences of exactly n words. *Hierarchy constraints* allow controlled generalization according to the item hierarchy to find patterns which do not directly occur in the input data. Examples include shopping patterns such as "customers frequently buy some *DSLR camera*, then some *tripod*, then some *flash*" or textual patterns such as "*PERSON be* born in *LOCATION*". *Regular expression (RE) constraints* have also been studied in the context of FSM; here subsequences must match a given RE.

A number of specialized algorithms for various combinations of the above subsequence constraints have been proposed in the literature. In this paper, we show that many subsequence constraints—including and beyond those described above—can be unified in a single framework. A unified treatment allows researchers to study subsequence constraints in general instead of focusing on certain combinations individually. It also helps to improve usability of pattern mining systems for practitioners because it avoids the need to develop customized mining algorithms for the particular subsequence constraint of interest. In this work, we focus on the questions of (1) how to model and express subsequence constraints in a suitable way and (2) how to mine efficiently all frequent sequences that satisfy the given constraints.

In more detail, we introduce *subsequence predicates* to model subsequence constraints in a general way, and we propose a simple and intuitive *pattern expression language* to concisely express subsequence predicates. Our pattern expressions are based on regular expressions, but—in contrast to prior work on RE-constrained FSM—target input sequences and support capture groups and item hierarchies. Capture groups are the key ingredient for expressing most prior subsequence constraints in a unified way; see Tab. I for examples. Direct support for item hierarchies allows us both to express subsequence constraints concisely and to mine generalized subsequences in a controlled way. Some example pattern expressions as well as anecdotal results are given in Tab. IV.

To mine frequent sequences, we propose to use finite state transducers (FST) as the underlying computational model. To the best of our knowledge, FSTs have not been studied in the context of FSM before. We propose the DESQ system, which includes two efficient mining algorithms termed DESQ-COUNT and DESQ-DFS. Both algorithms translate a given pattern expression to a *compressed* FST, which is subsequently optimized and simulated in a way suitable for frequent sequence mining. DESQ-COUNT is a match-and-count algorithm that aims at highly selective constraints, whereas DESQ-DFS can handle more demanding pattern expressions

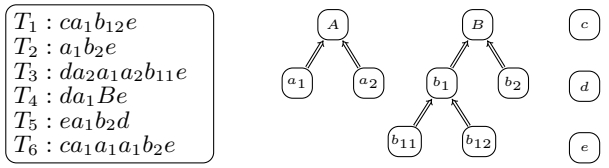

(a) Example database (b) Example hierarchy

Fig. 1: A sequence database and its vocabulary

and is inspired by PrefixSpan [11]. Our experimental study on various real-world datasets suggests that DESQ is an efficient general-purpose FSM framework and competitive to state-of-the-art specialized algorithms.

II. Preliminaries

Sequence database. A *sequence database* is a multiset of sequences, denoted $\mathscr{D} = \{ T_1, T_2, \ldots, T_{|\mathscr{D}|} \}$. Each *sequence* $T = t_1 t_2 \ldots t_{|T|}$ is an ordered list of items from a *vocabulary* $\Sigma = \{ w_1, w_2, \ldots, w_{|\Sigma|} \}$. We denote by ϵ the empty sequence, by $|T|$ the length of sequence T, by Σ^* (Σ^+) the set of all (all non-empty) sequences that can be constructed from items in Σ. Fig. 1(a) shows an example sequence database $\mathscr{D}_{ex}$ consisting of six sequences.

Item hierarchy. The items in Σ are arranged in an *item hierarchy*, which expresses how items can be generalized (or that they cannot be generalized). Fig. 1(b) shows an example hierarchy in which, for example, item a_1 generalizes to item A. In general, we say that an item u *directly generalizes to* an item v, denoted $u \Rightarrow v$, if u is a child of v. We further denote by $\Rightarrow^*$ the reflexive transitive closure of $\Rightarrow$. For the example of Fig. 1(b), we have $b_{11} \Rightarrow b_1$, $b_1 \Rightarrow B$ and $b_{11} \Rightarrow^* B$. For each item $w \in \Sigma$, we denote by $\mathrm{anc}(w) = \{ w' \mid w \Rightarrow^* w' \}$ the set of *ancestors* of w (including w) and by $\mathrm{desc}(w) = \{ w' \mid w' \Rightarrow^* w \}$ the set of *descendants* of w (again, including w). In our running example, we have $\mathrm{anc}(b_1) = \{ b_1, B \}$ and $\mathrm{desc}(b_1) = \{ b_1, b_{11}, b_{12} \}$.

Subsequence. Let $S = s_1 s_2 \ldots s_{|S|}$ and $T = t_1 t_2 \ldots t_{|T|}$ be two sequences composed of items from Σ. We say that S is a *generalized subsequence* of T, denoted $S \sqsubseteq T$, if S can be obtained by deleting and/or generalizing items in T. More formally, $S \sqsubseteq T$ iff there exists integers $1 \leq i_1 < i_2 < \cdots < i_{|S|} \leq |T|$ such that $t_{i_k} \Rightarrow^* s_k$ for $1 \leq k \leq |S|$. Continuing our example, we have $cBe \sqsubseteq T_1$, $ca_1 \sqsubseteq T_1$ and $a_1 c \not\sqsubseteq T_1$.

III. FSM With Subsequence Constraints

A subsequence constraint describes which subsequences of a given input sequence should be considered for frequent sequence mining. Our goal is to provide a general framework to express subsequence constraints, including and beyond previously proposed constraints. Consider the following (admittedly contrived) subsequence constraint as an example of a subsequence constraint.

Example 1. *Consider our example database $\mathscr{D}_{ex}$ and suppose that we are interested in mining sequences of B's and/or descendants of A's. We restrict attention to sequences that occur consecutively in input sequences starting with c or d and ending with e. We also allow to generalize occurrences of descendants of A and B. Then $a_1 B \sqsubseteq T_1$ and $AB \sqsubseteq T_1$ satisfies this subsequence constraint, whereas $a_1 b_{12} \sqsubseteq T_1$, $a_1 b_1 \sqsubseteq T_1$, $a_1 B \sqsubseteq T_2$ and $AB \sqsubseteq T_2$ do not.*

The above subsequence constraint cannot be expressed using prior methods. Note that the constraint combines (i) a gap constraint (consecutive), (ii) a hierarchy constraint (descendants of B must be generalized), and (iii) a context constraint (between c or d, and e).

Subsequence predicates. We propose subsequence predicates as a general, natural model for subsequence constraints. A *subsequence predicate* P is a predicate on pairs (S, T), where $T \in \Sigma^+$ is any input sequence and $S \sqsubseteq T$ is a subsequence. Subsequence $S \sqsubseteq T$ satisfies the constraint when $P(S, T)$ holds. Note that P is not a predicate on (only) subsequence S; it also involves input sequence T. We denote by $G_P(T) = \{ S \sqsubseteq T \mid P(S, T) \}$ the set of *P-subsequences* in T. For each $S \in G_P(T)$, we say that S is *P-generated* by T. For example, let P_{ex} be the subsequence predicate that expresses subsequence constraint of Ex. 1, then $G_{P_{ex}}(T_1) = \{a_1 B, AB\}$ and $G_{P_{ex}}(T_2) = \emptyset$.

Subsequence predicates can encode different application needs, including but not limited to the various subsequence constraints discussed before. A subsequence predicate can act as a filter on the set of all subsequences of T (only A's and B's), but may also consider the context in which these subsequences occur (consecutively between c or d and e). In practice, we may construct subsequence predicates that generate all n-grams, all adjective-noun pairs, all relational phrases between named entities, all electronic products, or, in log mining, sequences of items that occur before and/or after an error item. We propose a suitable way to express subsequence predicates in Sec. IV.

FSM and subsequence predicates. Let P be a subsequence predicate. The *P-support* $\mathrm{Sup}_P(S, \mathscr{D})$ of sequence $S \in \Sigma^+$ in sequence database $\mathscr{D}$ is the *multiset* of all sequences in $\mathscr{D}$ that P-generate S, i.e.,

$$\mathrm{Sup}_P(S, \mathscr{D}) = \{ T \in \mathscr{D} \mid S \in G_P(T) \}. \qquad (1)$$

The *P-frequency* of S in $\mathscr{D}$ is given by $f_P(S, \mathscr{D}) = |\mathrm{Sup}_P(S, \mathscr{D})|$. In our example database, we have $\mathrm{Sup}_{P_{ex}}(Aa_1AB, \mathscr{D}_{ex}) = \{T_3, T_6\}$ and thus $f_{P_{ex}}(Aa_1AB, \mathscr{D}_{ex}) = 2$. Given a *support threshold* $\sigma > 0$, we say that a sequence S is *P-frequent* if $f_P(S, \mathscr{D}) \geq \sigma$.

Problem Statement. *Given a sequence database $\mathscr{D}$, a subsequence predicate P, and a support threshold $\sigma > 0$, find all P-frequent sequences $S \in \Sigma^+$ along with their frequencies.*

The set of all P_{ex}-frequent sequences for $\sigma = 2$ in our example database is given by

$$\{ AAAB : 2, AB : 2, Aa_1AB : 2, a_1B : 2 \},$$

where we also give P-frequencies.

IV. Pattern Expressions

We propose a pattern language for expressing subsequence predicates in a simple and intuitive way. Our language is based on regular expressions, but adds features that allows us to unify many prior subsequence constraints. We subsequently suggest a computational model based on FSTs, and describe the formal semantics of our language.

A. Pattern Language

Our language consists of the following set of *pattern expressions*, defined inductively:

1) For each item $w \in \Sigma$, the expressions w, $w_{=}$, $w^{\uparrow}$, and $w^{\uparrow}_{=}$ are pattern expressions.
2) . and $.^{\uparrow}$ are pattern expressions.
3) If E is a pattern expression, so are (E), $[E]$, $[E]^*$, $[E]^+$, $[E]?$, and for all $n, m \in \mathbb{N}$ with $n \leq m$, $[E]\{n\}$, $[E]\{n,\}$, and $[E]\{n,m\}$.
4) If E_1 and E_2 are pattern expressions, so are $[E_1 E_2]$ and $[E_1|E_2]$.

Pattern expressions are based on regular expressions, but additionally include capture groups (in parentheses), hierarchies (by omitting $_{=}$), and generalizations (using $^{\uparrow}$ and $^{\uparrow}_{=}$). We make use of the usual precedence of rules for regular expressions to suppress square brackets (but not parentheses); operators that appear earlier in the above definition have higher precedence. We refer to expressions of form (1) or (2) as *item expressions*. We write $G_E(T)$ to refer to the set of subsequences "generated" by expression E on input T (see Sec. IV-B for a formal definition).

Captured and uncaptured expressions. Pattern expressions specify which subsequences to output (captured) as well as the context in which these subsequences should occur (uncaptured). We make use of parentheses to distinguish these two cases; the semantics is similar to the use of capture groups in regular expressions. Given an expression E, only subexpressions that are enclosed in or contain a capture group will produce non-empty output; all other subexpressions serve to describe context information. For example, the pattern expression

$$E_{ex} = [c|d]([A^{\uparrow} \mid B^{\uparrow}_{=}]^+)e. \tag{2}$$

describes precisely the subsequence constraint of Ex. 1. Here subexpressions $[c|d]$ and e describe context and $[A^{\uparrow} \mid B^{\uparrow}_{=}]^+$ output.

Item expressions. Item expressions are the elementary form of pattern expressions and apply to one input item. If the item expression "matches" the input item, it can "produce" an output item; see Tab. II for an overview. Fix some $w \in \Sigma$. The most basic item expression is $w_{=}$: it matches only item w and produces either ϵ (if uncaptured) or w (if captured). Using our example hierarchy of Fig. 2(a), we have $G_{A_{=}}(A) = \emptyset$ (note that we ignore output ϵ), $G_{(A_{=})}(A) = \{\, A \,\}$, and $G_{(A_{=})}(a_1) = \emptyset$. Sometimes we do not want to only match the specified item but also all of its descendants in the item hierarchy (e.g., we want to match all nouns in text mining). Item expression w serves this purpose: it matches any item $w' \in \mathrm{desc}(w)$ (which includes w) and, when captured, produces the item that has been matched. For example, we have $G_{(A)}(A) = \{\, A \,\}$, $G_{(A)}(a_1) = \{\, a_1 \,\}$, and $G_{(A)}(b_1) = \emptyset$. Our language also provides *wild card* symbol "." to match any item; again, the matched item is produced when the wild card is captured. For example, $G_{(.)}(A) = \{\, A \,\}$, and $G_{(.)}(a_1) = \{\, a_1 \,\}$.

To support mining with controlled generalizations (e.g., to mine patterns such as "*PERSON* lives in *CITY*"), we use the *generalization operator* $^{\uparrow}$, which generalizes items along the hierarchy. Item expressions that use the generalization operator must be captured. More specifically, item expression $w^{\uparrow}$ matches any item $w' \in \mathrm{desc}(w)$—as expression w does—, and it produces either the matched input item or any of its ancestors that is also a descendant of w. For example, $G_{(B^{\uparrow})}(b_{12}) = \{\, b_{12}, b_1, B \,\}$ and $G_{(b_1^{\uparrow})}(b_{12}) = \{\, b_{12}, b_1 \,\}$. We also allow the use of a wild card with generalization operator: expression ".$^{\uparrow}$" matches any item and produces each of its generalizations. For example, $G_{(.^{\uparrow})}(b_1) = \{\, b_1, B \,\}$. Our final item expression is used to enforce a generalization: $w^{\uparrow}_{=}$ matches any descendant of w and produces w, independently of which descendant has been matched. For example $G_{(B^{\uparrow}_{=})}(b_{12}) = \{\, B \,\}$.

TABLE I: Pattern expr. for prior subsequence constraints

Subsequence constraint	*Example*	Pattern expression
All subsequences [1], [11], [15]		$[.^*(.)]^+$
Bounded length [13]	*length 3–5*	$[.^*(.)]\{3,5\}$
n-grams [10], [14]	*3-, 4- and 5-grams*	$(.)\{3,5\}$
Bounded gap [10], [13]	*each gap at most 3*	$(.)[.\{0,3\}(.)]+$
Serial episodes [16]	*length 3, total gap* ≤ 2	$(.)[.?.?(.) \mid .?(.).? \mid (.).?.?](.)$
Hierarchy [1], [8]	*generalized 5-grams*	$(.^{\uparrow})\{5\}$
Regular expression [9], [12], [17], [18]		
	subsequences matching $[a\|b]\,c^*d$	$(a\|b)[.^*(c)]^*.^*(d)$
	contiguous subsequences matching $[a\|b]\,c^*d$	$([a\|b]\,c^*d)$

Composite expressions. Item expressions can be arbitrarily combined using operators ? (optionality), $*$ (Kleene star), $+$ (Kleene plus), $\{n,m\}$ (bounded repetition), $|$ (union), and concatenation to match (sequences of) more than one input item. The semantics of these compositions is as in regular expressions.

Examples. Our pattern expressions allow us to express many existing subsequence constraints in a unified way; see Tab. I for some examples. Note that the use of capture groups enables many of these pattern expressions. Pattern expressions can additionally express many customized subsequence constraints that cannot be handled by existing FSM frameworks; see Tab. IV for some examples.

B. Computational Model

We translate patterns expressions into FSTs, which are a natural computational model for pattern expressions. An FST is a type of finite state machine for string-to-string translation [19]. FSTs are similar to finite state automata but additionally label transitions with output strings. Conceptually, an FST reads an input string and translates it to an output string in a nondeterministic fashion. We will use FSTs to specify subsequence predicate $P(S,T)$: the predicate holds if the FST can output subsequence S when reading input T.

Finite state transducers. More formally, we consider a restricted form of FSTs defined as follows. An FST $\mathcal{A}$ is a 5-tuple $(Q, q_S, Q_F, \Sigma, \Delta)$, where Q is a set of states, $q_S \in Q$ is the initial state, $Q_F \subseteq Q$ is the set of final states, Σ is an input and output alphabet, and $\Delta \subseteq Q \times (\Sigma \cup \{\,\epsilon\,\}) \times (\Sigma \cup \{\,\epsilon\,\}) \times Q$ is a transition relation. For every transition $(q_{from}, in, out, q_{to}) \in \Delta$, we require that $out \in \mathrm{anc}(in) \cup \{\,\epsilon\,\}$ and that whenever $in = \epsilon$ then $out = \epsilon$. Our notion of FSTs differs from traditional FSTs in that we use a common input and output alphabet and in that we restrict output labels. The latter restriction ensures that our FSTs output generalized subsequences of their input (Lemma 1). Fig. 2(a) shows an example FST, where $q_S = q_0$, $Q_F = \{\, q_{11} \,\}$, and each transition is marked with in:out labels. We refer to transitions with $in = \epsilon$ (and thus $out = \epsilon$) as *ϵ-transitions*; these transitions are marked with ϵ in the figure.

Runs and outputs. Let $T = t_1 t_2 \ldots t_n$ be an input sequence. A *run* for T is a sequence $p = p_1 p_2 \ldots p_m$ of transitions, where for $1 \leq i \leq m$: $p_i = (q_i, w_i, w'_i, q'_i) \in \Delta$, $q_1 = q_S$,

TABLE II: Translation rules for item expressions (where $w, w', w'' \in \Sigma$)

Expr.	Matches	Transl. type	Produces	FST	Compressed FST
$w_=$	w	Uncaptured	ϵ	$\{q_S \xrightarrow{w:\epsilon} q_F\}$	$\{q_S \xrightarrow{w_=:\epsilon} q_F\}$
		Captured	w	$\{q_S \xrightarrow{w:w} q_F\}$	$\{q_S \xrightarrow{w_=:w} q_F\}$
w	$w' \in \text{desc}(w)$	Uncaptured	ϵ	$\{q_S \xrightarrow{w':\epsilon} q_F \mid w' \in \text{desc}(w)\}$	$\{q_S \xrightarrow{w:\epsilon} q_F\}$
		Captured	w'	$\{q_S \xrightarrow{w':w'} q_F \mid w' \in \text{desc}(w)\}$	$\{q_S \xrightarrow{w:\$} q_F\}$
.	$w \in \Sigma$	Uncaptured	ϵ	$\{q_S \xrightarrow{w:\epsilon} q_F \mid w \in \Sigma\}$	$\{q_S \xrightarrow{.:\epsilon} q_F\}$
		Captured	w	$\{q_S \xrightarrow{w:w} q_F \mid w \in \Sigma\}$	$\{q_S \xrightarrow{.:\$} q_F\}$
$w^\uparrow$	$w' \in \text{desc}(w)$	Captured	$\text{anc}(w') \cap \text{desc}(w)$	$\{q_S \xrightarrow{w':w''} q_F \mid w' \in \text{desc}(w), w'' \in \text{anc}(w') \cap \text{desc}(w)\}$	$\{q_S \xrightarrow{w:\$\text{-}w} q_F\}$
$.^\uparrow$	$w \in \Sigma$	Captured	$\text{anc}(w)$	$\{q_S \xrightarrow{w:w'} q_F \mid w \in \Sigma, w' \in \text{anc}(w)\}$	$\{q_S \xrightarrow{.:\$\text{-}\top} q_F\}$
$w_=^\uparrow$	$w' \in \text{desc}(w)$	Captured	w	$\{q_S \xrightarrow{w':w} q_F \mid w' \in \text{desc}(w)\}$	$\{q_S \xrightarrow{w:w} q_F\}$

$q_{i+1} = q'_i$, and $w_1 w_2 \dots w_m = T$ (recall that $w_i \in \Sigma \cup \{\epsilon\}$ so that $m \geq n$). Intuitively, the FST starts in state q_S and repeatedly selects transitions that are consistent with the next input item. If $q_m \in Q_F$, we refer to p as an *accepting run*. The output $O(p)$ of run p is the sequence $S = w'_1 \dots w'_m$ of output labels, where we omit all w'_i with $w'_i = \epsilon$ and set $S = \epsilon$ if all $w'_i = \epsilon$. The set of sequences generated by FST $\mathcal{A}$ is given by

$$G_\mathcal{A}(T) = \{\, O(p) \neq \epsilon \mid p \text{ is an accepting run of } \mathcal{A} \text{ for } T \,\}.$$

EXAMPLE 2. *Consider the FST $\mathcal{A}_{F2(a)}$ of Fig. 2(a). $\mathcal{A}_{F2(a)}$ has two accepting runs for sequence $T_1 = ca_1b_{12}e$, which are given by $p_1 = q_0 \xrightarrow{\epsilon} q_1 \xrightarrow{c:\epsilon} q_3 \xrightarrow{\epsilon} q_5 \xrightarrow{\epsilon} q_6 \xrightarrow{a_1:a_1} q_8 \xrightarrow{\epsilon} q_{10} \xrightarrow{\epsilon} q_5 \xrightarrow{\epsilon} q_7 \xrightarrow{b_{12}:B} q_9 \xrightarrow{\epsilon} q_{10} \xrightarrow{e:\epsilon} q_{11}$ with output $O(p_1) = a_1B$, and $p_2 = q_0 \xrightarrow{\epsilon} q_1 \xrightarrow{c:\epsilon} q_3 \xrightarrow{\epsilon} q_5 \xrightarrow{\epsilon} q_6 \xrightarrow{a_1:A} q_8 \xrightarrow{\epsilon} q_{10} \xrightarrow{\epsilon} q_5 \xrightarrow{\epsilon} q_7 \xrightarrow{b_{12}:B} q_9 \xrightarrow{\epsilon} q_{10} \xrightarrow{e:\epsilon} q_{11}$. Thus, $G_{\mathcal{A}_{F2(a)}}(T_1) = \{\, a_1B, AB \,\}$, as desired. There is no accepting run for T_2 so that $G_{\mathcal{A}_{F2(a)}}(T_2) = \emptyset$. Observe that $\mathcal{A}_{F2(a)}$ generates precisely the P-sequences of Ex. 1.*

The following lemma states that our FSTs generate generalized subsequences of their inputs and thus specify subsequence predicates. Note that the lemma holds for any run, whether or not accepting.

LEMMA 1. *Let $T \in \Sigma^*$ be an input sequence and $\mathcal{A}$ be an FST. For any run p of $\mathcal{A}$ for T, it holds $O(p) \sqsubseteq T$.*

Note that not all subsequence predicates can be expressed with FSTs; e.g., there is no FST for predicate "all subsequences of form a^*b^* with an equal number of a's and b's". FST are a good trade-off between expressiveness and computational complexity, however: they can express many subsequence constraints that occur in practice and they lend themselves to efficient mining (see Sec. V).

Translating pattern expression. We now describe how to translate a pattern expression E into an FST $\mathcal{A}(E)$. The FST formally defines the semantics of pattern expressions: we set $G_E(T) \stackrel{\text{def}}{=} G_{\mathcal{A}(E)}(T)$. Each item expression is translated into a two-state FST with $Q = \{\, q_S, q_F \,\}$, where q_S is the initial and q_F the final state. The transitions of the FST depend on the item expression and are summarized in Tab. II, column "FST". The translation rules for composite expressions mirror the Thompson construction [20] for translating regular expressions to finite state automata. [1]For example, expression E_{ex} of Eq. (2) translates to the FST of Fig. 2(a).

Compressed FST. The translation rules above can produce very large FSTs, especially when the vocabulary is large. For example, if the hierarchy has n items and average depth d, the FST for ".$^\uparrow$" has $\Theta(nd)$ transitions. To avoid this explosion of FST size and support efficient mining, we make use of a compressed FST (cFST) representation for this purpose; see column "compressed FST" of Tab. II. The cFST of an item expression has exactly one transition, but input and output labels are taken from an alphabet larger than Σ. Each transition in the cFST describes a set of transitions in the corresponding FST in a concise way. More specifically, cFSTs use as input labels ., w, and $w_=$ for all $w \in \Sigma$. Here "." matches all input items, w matches all items in $\text{desc}(w)$, and $w_=$ matches only item w. cFSTs use as output labels ϵ, w, \$, \$-w, and \$-$\top$ for $w \in \Sigma$. Each transition encodes the set of output labels in the corresponding FST: ϵ and w are as before, \$ encodes the matched input item, \$-w the matched input item and all its ancestors up to w, and \$-$\top$ the matched item and all its ancestors. The cFST translations for composite expressions remain unmodified. Fig. 2(b) shows a cFST $\mathcal{A}_{ex}$ for E_{ex}, Eq. (2). Note that the cFST has fewer transitions than its uncompressed counterpart of Fig. 2(a).

Simulating cFSTs. A simple way to compute the set $G_\mathcal{A}(T)$ for a cFST is via simulation and backtracking. Note that the computation of $G_\mathcal{A}(T)$ for all $T \in \mathcal{D}$ can be infeasible. Nevertheless, simulation forms the basis of the more efficient DESQ-DFS algorithm of Sec. V-C so that we describe the approach briefly. Denote by the *transition function*

$$\delta(q, w) = \{\, (out, q_{to}) \mid (q, in, out, q_{to}) \in \Delta,\ in \text{ matches } w \,\}$$

the set of (output label, state)-pairs that can be reached from state q by consuming input item w (see column "Matches" in Tab. II). To simulate a cFST, we start with the initial state q_S and repeatedly select a transition from $\delta(q, w)$, where q is the current state and w the next input item. If there are multiple such transitions (i.e., when $|\delta(q, w)| > 1$) , we try them one by one via backtracking. As we move from state to state, we keep track of the outputs in a buffer (column "Produces" in Tab. II). If we reach a final state after consuming all input items, we add the buffered output to the set $G_\mathcal{A}(T)$.

Partial matches. The simulation algorithm only generates an output when the entire input sequence is matched. If we are interested in matching pattern expressions that occur somewhere in the input sequence instead, we construct a cFST for $.^*E$ (instead of for E) and modify the above simulation such that it adds the buffered output to $G_\mathcal{A}(T)$ whenever a

[1]All translation rules can be implemented without introducing any ϵ-transitions; we follow this approach in our actual implementation but use ϵ-transitions in our example FSTs for improved readability.

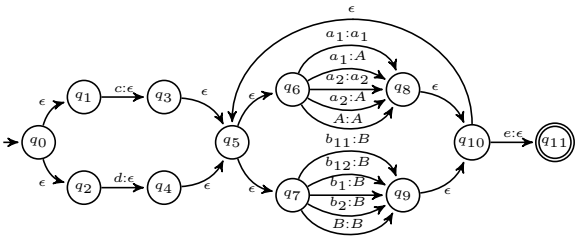

(a) FST

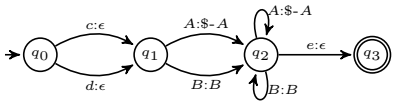

(b) Compressed FST

Fig. 2: FST (a) and cFST (b) for $[c|d]([A^{\uparrow} \mid B^{\uparrow}_{=}]^{+})e$.

final state is reached, whether or not the entire input has been consumed.

Nondeterminism. Note that cFST simulation involves backtracking when multiple transitions match the same input item and/or a transition has an output label of form \$-w or \$-$\top$. The standard way to avoid non-determinism is to use some form of FST determinization [19]. In general, these methods do not directly apply to our FSTs because there are no "sequential" or even "p-subsequential" transducers for some pattern expressions (e.g. $E = [.^{*}(.)]^{+}$). In our implementation, we adapt the power construction algorithm [21] and some other heuristics to reduce non-determinism to the extent possible. The remaining non-determinism (if any) did not lead to a bottleneck in our experimental study.

V. Pattern Mining

We now turn attention to mining P-frequent sequences from a sequence database. We assume that subsequence predicate P is described by a cFST $\mathcal{A}$ (e.g., obtained by translating a pattern expression). We propose three methods for mining P-frequent sequences: Naïve, DESQ-COUNT, and DESQ-DFS.

A. Naïve Approach

The naïve "generate-and-count" approach is to compute $G_{\mathcal{A}}(T)$ for each input sequence $T \in \mathscr{D}$ via cFST simulation and count how often each sequence has been generated (cf. Eq. (1)). The naïve approach is generally inefficient because it considers many globally infrequent sequences. For example, we obtain

$$G_{A_{ex}}(T_3) = \{AAAB, AAa_2B, Aa_1AB, Aa_1a_2B, a_2AAB, a_2Aa_2B, a_2a_1AB, a_2a_1a_2B\}$$

for input sequence T_3, but only $AAAB$ and Aa_1AB are actually P-frequent.

B. DESQ-COUNT

DESQ-COUNT reduces the number of sequences that are generated and counted by making use of item frequencies. In more detail, denote by $f(w, \mathscr{D}) = |\{T \in \mathscr{D} \mid w \sqsubseteq T\}|$ the *frequency* of item w. We say that item w is *frequent* if $f(w, \mathscr{D}) \geq \sigma$. Similar to many prior FSM algorithms, DESQ-COUNT first generates an *f-list* F, which contains all items along with their frequencies. For our example database, we obtain f-list

$$F_{ex} = \{A{:}6, e{:}6, B{:}6, a_1{:}6, d{:}3, b_2{:}3, b_1{:}2, c{:}2, b_{12}{:}1, b_{11}{:}1, a_2{:}1\}. \quad (3)$$

Note that the f-list is independent of the subsequence constraint and can be precomputed. In DESQ-COUNT, we make use of the f-list to reduce the size of $G_{\mathcal{A}}(T)$. Denote by

$$G^{F}_{\mathcal{A}}(T) = \{\, S \in G_{\mathcal{A}}(T) \mid \forall w \in S : f(w, \mathscr{D}) \geq \sigma \,\}$$

the subset of generated sequences that do not contain infrequent items. For T_3, we have $G^{F_{ex}}_{\mathcal{A}_{ex}}(T_3) = \{AAAB, Aa_1AB\}$, which is much smaller than the full set $G_{\mathcal{A}_{ex}}(T_3)$ given above. DESQ-COUNT proceeds as the naïve approach, but replaces $G_{\mathcal{A}}(T)$ by $G^{F}_{\mathcal{A}}(T)$ for each $T \in \mathscr{D}$. Note that we do not fully compute $G_{\mathcal{A}}(T)$ to obtain $G^{F}_{\mathcal{A}}(T)$, instead we directly compute the reduced set $G^{F}_{\mathcal{A}}(T)$ by adapting cFST simulation to work with the f-list. To do so, we stop exploring a run as soon as an infrequent item is produced.

The correctness of DESQ-COUNT is established by Lemma 1, which states that FSTs specify subsequence predicates, and the following observation.

Lemma 2. *Let P be a subsequence predicate and $S \in \Sigma^{+}$ be any sequence. Then for all $w \in S$, $f(w, \mathscr{D}) \geq f_P(S, \mathscr{D})$.*

Proof: Pick any $w \in S$ and input sequence $T \in \mathscr{D}$ such that $S \in G_P(T)$. Since P is a subsequence predicate, $S \sqsubseteq T$. Since $w \in S$, we have $w \sqsubseteq S$ and thus also $w \sqsubseteq T$. We obtain

$$\begin{aligned} f_P(S, \mathscr{D}) &= |\{\, T \in \mathscr{D} \mid S \in G_P(T) \,\}| \\ &\leq |\{\, T \in \mathscr{D} \mid w \sqsubseteq T \,\}| = f(w, \mathscr{D}) \end{aligned}$$ ■

The pruning performed by DESQ-COUNT can substantially reduce the number of candidate sequences. DESQ-COUNT is inefficient (and sometimes infeasible), however, if pruning is not sufficiently effective and the sets $G^{F}_{\mathcal{A}}(T)$ are very large. The DESQ-DFS algorithm, which we present next, addresses such cases.

C. DESQ-DFS

DESQ-DFS adapts the pattern-growth framework of PrefixSpan [11] to FSTs. Pattern growth approaches arrange the output sequences in a tree, in which each node corresponds to a sequence S and is associated with a *projected database*, which consists of the set of input sequences in which S occurs. Starting with an empty sequence and the full sequence database, the tree is built recursively by a performing series of *expansions*. In each expansion, a frequent sequence S (of l items) is expanded to generate sequences with prefix S (of $l{+}1$ items), their projected databases, and their supports. In what follows, we describe how we adapt these concepts to mine P-frequent sequences; the corresponding algorithm for cFSTs is shown as Alg. 1 and illustrated on our running example in Fig. 3.

Projected databases. For a sequence S, we store in its projected database the state of the simulations of $\mathcal{A}$ on all input sequences that generate S as a partial output. We refer to such a state as a *snapshot* for S. The snapshot concisely describes which items have been consumed, which state the FST simulation is in, and which output has been produced so far. In more detail, suppose that we simulate an $\mathcal{A}$ on input sequence $T = t_1 \dots t_n$. Consider a partial run $p = p_1 \dots p_m$

Algorithm 1 DESQ-DFS

Require: $\mathscr{D}$, cFST $\mathcal{A} = (Q, q_S, Q_F, \Sigma, \Delta)$, σ, f-list F
Ensure: P-frequent sequences for $\mathcal{A}$ in $\mathscr{D}$

```
 1: S ← ϵ                    // create root node; initially fields S.Proj = S.Sup = ∅
 2: S.Proj ← { T_1[1@q_S], ..., T_|D|[1@q_S] }
 3: EXPAND(S)
 4:
 5: void EXPAND(S):
 6: for all T[pos@q] ∈ S.Proj do        // simulate cFST for all snapshots
 7:     INCSTEP(T, pos, q, S)
 8: if |S.Sup| ≥ σ then yield (S,|S.Sup|)        // Output if P-frequent
 9: for all S' ∈ S.Children do    // expand if prefix support large enough
10:     if |{ T | ∃pos, q: T[pos@q] ∈ S.Proj }| ≥ σ then EXPAND(S')
11:
12: void INCSTEP(T, pos, q, S):        // simulate until an item is produced
13: if q ∈ Q_F and pos > |T| and S ≠ ϵ then
14:     S.Sup ← S.Sup ∪ { T }                                  // initially empty
15: for all (out, q_to) ∈ δ(q, t_pos) do
16:     switch (out)
17:         case ϵ:
18:         INCSTEP(T, pos+1, q_to, S)
19:         case w:
20:         if f(w, D) ≥ σ then APPEND(S, w, T, pos+1, q_to)
21:         case $:
22:         if f(t_pos, D) ≥ σ then APPEND(S, t_pos, T, pos+1, q_to)
23:         case $-x, x ∈ Σ ∪ { ⊤ }:
24:         for all w' ∈ anc(t_pos) ∩ desc(x) do
25:             if f(w', D) ≥ σ then APPEND(S, w', T, pos+1, q_to)
26:
27: void APPEND(S, w, T, pos, q):
28: S.Children ← S.Children ∪ { Sw }      // node Sw is created if new
29: Sw.Proj ← Sw.Proj ∪ { T[pos@q] }                       // initially empty
```

obtained after $m \leq n$ steps. We generated output $S = O(p)$ and, under our running assumption that $\mathcal{A}$ does not contain ϵ-transitions, consumed prefix $T' = t_1 \dots t_m$ of T at this time. If the output item of the last transition p_m is not empty (and thus agrees with the last item of S), we say that triple $T[pos@q]$ is a *snapshot* for S, where $pos = m + 1$ is the position of next input item and q is the last state in p (current state of $\mathcal{A}$). The *projected database* for S consists of all snapshots for S and is given by

$$\mathrm{Proj}_{\mathcal{A}}(S, \mathscr{D}) = \{\, T[pos@q] \mid T \in \mathscr{D}, T[pos@q] \text{ is a snapshot for } S \text{ on } \mathcal{A} \,\}.$$

Fig. 3(b) shows some projected databases associated with some sequences for our running example. For example, we obtained the partial output a_1 only from input sequences T_1, T_4, and T_6. In each case, we consumed two items (next item is at position 3) and ended in state q_2. We refer to the number of input sequences that can generate S as a partial output as the *prefix support* of S:

$$\mathrm{Presup}_{\mathcal{A}}(S, \mathscr{D}) = \{T \mid \exists pos, q : T[pos@q] \in \mathrm{Proj}_{\mathcal{A}}(S, \mathscr{D})\}.$$

In our example, $\mathrm{Presup}_{\mathcal{A}_{F3}}(a_1, \mathscr{D}_{ex}) = \{\, T_1, T_4, T_6 \,\}$. Note that even if an input sequence has multiple snapshots for S, it contributes only once to the prefix support.

Expansions. We start with root node ϵ and all snapshots for ϵ (lines 1 and 2) and then perform a series of expansions (lines 3 and 10). In each expansion, we scan the projected database sequentially. For each snapshot $T[pos@q]$ (lines 6–7), we resume the FST for T at item t_{pos} in state q (via INCSTEP, lines 12–25). The transducer is stopped as soon as an output item is produced or the entire input is consumed. In the former case, suppose we produce item w after consuming k more input items from T and thereby reach state q'. We then add the new snapshot $T[pos{+}k@q']$ to the projected database of child node Sw (lines 20, 22, and 25). In the later case, if

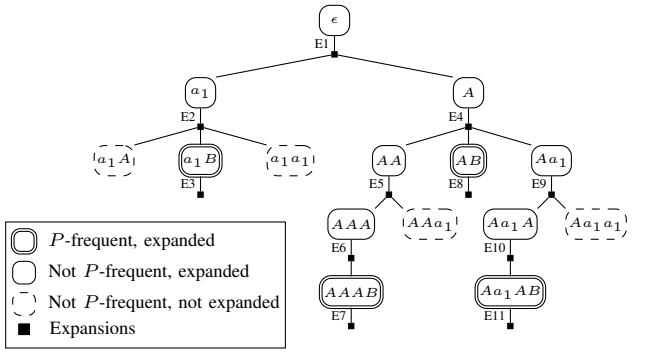

(a) Search space

S	S.Proj	$\lvert S.\mathrm{Presup}\rvert$	$\lvert S.\mathrm{Sup}\rvert$
ϵ	$\langle T_1[1@q_0], T_2[1@q_0], T_3[1@q_0], T_4[1@q_0], T_5[1@q_0], T_6[1@q_0]\rangle$	6	0
a_1	$\langle T_1[3@q_2], T_4[3@q_2], T_6[3@q_2]\rangle$	3	0
a_1A	$\langle T_6[4@q_2]\rangle$	1	0
a_1B	$\langle T_1[3@q_2], T_4[3@q_2]\rangle$	2	2
a_1a_1	$\langle T_6[4@q_2]\rangle$	1	0

(b) Some projected databases, prefix supports, and supports

Fig. 3: Illustration of DESQ-DFS for $\mathscr{D}_{ex}$, $\mathcal{A}_{Fex}$, and $\sigma = 2$

we end up in a final state (lines 13–14), we conclude that $T \in \mathrm{Sup}_{\mathcal{A}}(S, \mathscr{D})$ (see below). For example, both snapshots of a_1B reach final state q_3 by consuming all input items and without producing further output, so that $a_1B.\mathrm{Sup} = \{\, T_1, T_4 \,\}$.

Pruning. The above expansion procedure allows us to prune partial sequences as soon as it becomes clear that they cannot be expanded to a P-frequent sequence. We use two pruning techniques. First, as in DESQ- COUNT, we consider item w only if it is frequent; otherwise, we ignore the new snapshot. For example, when expanding a_1, we do not create nodes for sequences that contain infrequent items; e.g., a_1b_{12} has snapshot $T_1[4@q_2]$ but contains infrequent item b_{12} (see Eq. (3)). Second, we expand only those nodes S that have a sufficiently large prefix support—i.e., $\mathrm{Presup}_{\mathcal{A}}(S, \mathscr{D}) \geq \sigma$—and stop as soon as there is no such node anymore. For example, we do not expand node a_1a_1 because it contains only one snapshot, but we require snapshots from at least $\sigma = 2$ different input sequences.

Correctness. Note that the size of the prefix support is monotonically decreasing as we go down the tree but always stays at least as large as the support. This property, which we establish next, is key to the correctness of DESQ-DFS.

LEMMA 3. *For any sequence $S \in \Sigma^*$ and item $w \in \Sigma$, we have* $\mathrm{Presup}_{\mathcal{A}}(Sw, \mathscr{D}) \subseteq \mathrm{Presup}_{\mathcal{A}}(S, \mathscr{D})$.

Proof: Pick any $S \in \Sigma^*$, $w \in \Sigma$, and $T = t_1 \dots t_n \in \mathscr{D}$ with $T \in \mathrm{Presup}_{\mathcal{A}}(Sw, \mathscr{D})$. Then there exists a run $p = p_1 \dots p_m$ for prefix $T' = t_1 \dots t_m$ and some $m \leq n$ such that $O(p) = Sw$. Recall that inputs (outputs) are consumed (generated) from left to right. We conclude that there exists some $m' < m$ such that run $p' = p_1 \dots p_{m'}$ satisfies $O(p') = S$. Pick the shortest such run; then $p_{m'}$ outputs the last item of S. Since p' is additionally a run for $t_1 \dots t_{m'}$, which is a prefix of T, we conclude that $T \in \mathrm{Presup}_{\mathcal{A}}(S, \mathscr{D})$. ■

We now establish the correctness of DESQ-DFS.

THEOREM 1. *DESQ-DFS outputs each P-frequent sequence $S \in \Sigma^+$ with frequency $f_P(S, \mathscr{D})$. No other sequences*

are output.

Proof: Let $\mathcal{A} = (Q, q_S, Q_F, \Sigma, \Delta)$ be an FST and pick any sequence $S \in \Sigma^+$. We start with showing that Alg. 1 correctly computes the P-support of S when expanding node S, i.e., $S.\mathrm{Sup} = \mathrm{Sup}_{\mathcal{A}}(S, \mathscr{D})$ after the expansion. First pick any $T \in \mathrm{Sup}(S, \mathscr{D})$ with $T = t_1 \dots t_n$. Then there is an accepting run $p = p_1 \dots p_n$ for T. By arguments as in the proof of Lemma 3, there must be a smallest run $p' = p_1 \dots p_m$, $m \leq n$, such that $O(p') = S$ as well. Let q_m (q_n) be the state reached in transition p_m (p_n). We conclude that snapshot $T[pos@q_m] \in \mathrm{Proj}_{\mathcal{A}}(S, \mathscr{D})$, where $pos = m + 1$, and thus $T \in \mathrm{Presup}(S, \mathscr{D})$. Since by definition $p_{m+1} \dots p_n$ must output ϵ, Alg. 1 follows transitions $p_{m+1} \dots p_n$ without stopping when resuming snapshot $T[pos@q_m]$. By doing so, it consumes all the remaining items $t_{m+1} \dots t_n$ of T and reaches final state q_n. It thus includes T into $S.\mathrm{Sup}$ (lines 13–14). Now pick $T \notin \mathrm{Sup}_{\mathcal{A}}(S, \mathscr{D})$. Since there is no accepting run for T, Alg. 1 cannot reach a final state after consuming T so that it does not include T into $S.\mathrm{Sup}$. Putting both together, $S.\mathrm{Sup} = \mathrm{Sup}_{\mathcal{A}}(S, \mathscr{D})$ after expanding S, as desired. We conclude that Alg. 1 computes the correct frequency $f_P(S, \mathscr{D}) = |\mathrm{Sup}_{\mathcal{A}}(S, \mathscr{D})|$. S is output only if it is P-frequent (line 8). Note that for $S = \epsilon$, we have $\epsilon.\mathrm{Sup} = \emptyset$ (see line 13) so that ϵ is not output.

Let $S \in \Sigma^+$ be a P-frequent sequence. It remains to show that Alg. 1 reaches and expands node S. First observe that for any prefix S' of S, we have

$$\mathrm{Presup}(S', \mathscr{D}) \supseteq \mathrm{Presup}(S, \mathscr{D}) \supseteq \mathrm{Sup}(S, \mathscr{D}).$$

Here the first inclusion follows from Lemma 3, and the second inclusion follows from the above arguments. Since S is P-frequent, we have $|\mathrm{Sup}(S, \mathscr{D})| \geq \sigma$, which implies $|\mathrm{Presup}(S', \mathscr{D})| \geq \sigma$. Since every node on the path from ϵ to S corresponds to a prefix of S, Alg. 1 does not prune any of these nodes due to too low prefix support (line 10). To complete the proof, recall that S cannot contain an infrequent item by Lemma 2. Thus none of the nodes on the path from ϵ to S are pruned due to too low item frequency either (lines 20, 22, or 25). We conclude that Alg. 1 reaches and expands node S. ■

VI. Experimental Evaluation

We conducted an experimental study on three publicly available real-world datasets: a collection of text documents (for text mining), a collection of product reviews (for customer behavior mining), and a collection of protein sequences. Our goal was to investigate whether pattern expressions are sufficiently powerful to express prior and new subsequence constraints, whether DESQ's algorithms are efficient, and how they perform relative to each other and to prior algorithms.

Summary of our results. (1) Many subsequence constraints can be expressed with pattern expressions. (2) DESQ has acceptable overhead over state-of-the-art specialized sequence miners for common subsequence constraints. (3) DESQ-COUNT was consistently faster than Naïve. (4) DESQ-COUNT and DESQ-DFS had similar performance in cases where the average number of P-subsequences per input sequence was small. (5) When many subsequences per input are generated, DESQ-DFS was more than an order of magnitude faster than DESQ-COUNT and Naïve. (6) cFSTs sped up pattern matching by multiple orders of magnitude when compared to the state-of-the-art FST library OpenFST. Our results indicate that DESQ is a suitable general-purpose system for a wide range of subsequence constraints.

TABLE III: Dataset statistics

		NYT	AMZN	PRT
Sequence database	# Sequences	21,590,967	6,643,666	103,120
	Avg. length	19.9	4.5	482
	Max. length	5,042	25,630	600
	Total items	430,279,662	29,667,966	49,729,890
	Distinct items	3,975,859	2,374,096	25
Hierarchy	Total items	4,136,774	2,385,775	103,120
	Leaf items	3,901,118	2,371,522	103,120
	Interm. items	235,633	11,630	0
	Root items	23	2,623	103,120
	Max. depth	3	8	1
	Avg. fan-out	17.5	204	0
	Max. fan-out	1,505,913	332,723	0

A. Experimental Setup

Datasets. Tab. III summarizes our datasets. NYT is a subset of the New York Times corpus [22] and contains news articles. We generated an item hierarchy using annotations from the Stanford CoreNLP tools. The NYT hierarchy consists of named entities, which generalize to their type (PERSON, ORGANIZATION, LOCATION, MISC) and then to ENTITY, and of words, which generalize to their lemma and then to their part-of-speech tag. For example, "Maradona"⇒PERSON ⇒ENTITY and "is"⇒"be"⇒VERB.

AMZN is a dataset of Amazon product reviews [23] from which we extracted sequences of products (ordered by review timestamps) for each user. We used the Amazon product hierarchy as our item hierarchy. For example, "Canon 5D" ⇒"Digital Cameras"⇒"Camera & Photo"⇒ "Electronics".

PRT is a dataset of protein sequences [24] composed of 25 amino acid codes (items). The hierarchy is flat, i.e., there are no generalizations.

Implementation and setup. We implemented DESQ[2] in Java (JDK 1.8). We used ANTLR4 to generate a parser for pattern expressions. The cFST is constructed from the resulting parse tree. To measure the overhead of DESQ for common subsequence constraints, we compared it against state-of-the-art methods. For length and gap constraints, we used (1) C++ implementation of cSPADE [13] from the author, (2) our implementation of SPADE in Java that additionally handles hierarchy constraints, (3) our implementation of prefix-growth [12] in Java. For RE constraints, we used (1) prefix-growth and a C++ executable of SMA [18] obtained from the authors. To evaluate cFSTs we compared it against state-of-the-art FST library OpenFST 1.5.0 [25].

We preprocessed the datasets to compute the f-list and assign integer identifiers to each item. Item identifiers were assigned in descending order of item frequency, thus a more frequent item received a smaller item identifier. In our implementations, we encoded the sequence database compactly as arrays of item identifiers and use variable-length byte encoding to compress projected databases. Experiments on the NYT and AMZN datasets were performed on a machine with two Intel(R) Xeon(R) CPU E5-2640 v2 processors and 128GB of RAM running CentOS Linux 7.1. Experiments on the PRT dataset were performed on a machine equipped with Intel Core

[2]http://dws.informatik.uni-mannheim.de/en/resources/software/desq/

TABLE IV: Example pattern expressions for IE and NLP applications (N_1–N_5), customer behavior mining applications (A_1–A_4), protein sequence mining (P_1–P_4), and traditional sequence mining (T_1–T_3)

Pattern expression	Description	Example patterns from NYT dataset (frequency)
T_1: (.){1,λ}	*n-grams of up to λ words*	green tea (337), editor in chief (3275)
T_2: (.)[.{0,γ}(.)]{1,λ − 1}	*Skip n-grams with gap at most γ words and of up to length λ*	flight from to (758), son of and of (15896)
T_3: (.↑){1,λ}	*Generalized n-grams of up to λ words*	NOUN PREP DET NOUN (4.2M), PERSON be NOUN (2199)
		Example patterns from PRT dataset (frequency)
P_1: ([S|T]).*(.).*([R|K])	*subsequences that match RE≡[S|T].[R|K]*	SLR (103,093), TAK (102,941), SAK (102,946)
P_2: ([I|V]).*(D).*(L).*(G).*(T).*([S|T]).*(.).*([S|C])	*subsequences that match RE≡[I|V]DLGT[S|T].[S|C]*	IDLGTTLS (102,975), VDLGTSTC (92,662) VDLGTSDS (102,901)
P_3: ([S|T].[R|K])	*contiguous subsequences that match RE≡[S|T].[R|K]*	SLR (14,995), TAK (8,840), SAK (10,397)
P_4: ([S|T]..[D|E])	*contiguous subsequences that match RE≡[S|T]..[D|E]*	SDLE (2,015), TLEE (2,329), SGLD (1,054)
		Example patterns from NYT dataset (frequency)
N_1: ENTITY (VERB+ NOUN+? PREP?) ENTITY	*Relational phrase between entities*	lives in (847), is being advised by (15), has coached (10)
N_2: (ENTITY↑ VERB+ NOUN+? PREP? ENTITY↑)	*Typed relational phrases*	ORG headed by ENTITY (275), PER born in LOC (481)
N_3: (ENTITY↑ be=↑) DET? (ADV? ADJ? NOUN)	*Copular relation for an entity*	PER be novelist (165), LOC be great place (38),
N_4: (.↑){3} NOUN	*Generalized 3-grams before a noun*	NOUN PREP DET (4,223,219), DET ADV ADJ (350,005)
N_5: ([.↑ . .]|[. .↑ .]|[. . .↑])	*Generalized 3-grams, where at most one item is generalized*	the ADJ human (1,238), for DET book (1,704)
		Example patterns from AMZN dataset (frequency)
A_1: (Electr↑)[.{0,2}(Electr↑)]{1,4}	*Generalized sequences of (up to 5) electronic items, which are at most 2 items apart in the input sequences*	"Mice", "Keyboards", "Computers & Accessories" (556), "MP3 Players", "Headphones" (814)
A_2: (Book)[.{0,2}(Book)]{1,4}	*Sequences of books*	"The Bourne Supremacy", "The Bourne Ultimatum" (16)
A_3: DigitalCamera[.{0,3}(.↑)]{1,4}	*Type of products bought after a digital camera*	"Lenses", "Tripods" (11), "Batteries", "SD Cards" (12)
A_4: (MInstr↑)[.{0,2}(MInstr↑)]{1,4}	*Generalized sequences of musical instruments*	"Saxophones", "Bags & Cases", "Instr. Accessories" (127)

i7-4712HQ and 16GB RAM running Windows 10. We used a different setup for the PRT dataset as the SMA implementation is provided as a Windows binary only. All experiments were run single-threaded.

Methodology. For each experiment, we report the performance in terms of the total wall-clock time between launching the mining task and receiving the final result (including I/O). All measurements were averaged over three independent runs. Unless stated otherwise, all methods produced the same result.

B. Overhead of DESQ for Common Subsequence Constraints

We first investigated the overhead of DESQ compared to specialized miners for prior subsequence constraints. In particular we considered traditional as well as RE constraints.

Traditional constraints. We considered length and gap constraints as well as item hierarchies. We map these constraints to pattern expressions and obtain T_1–T_3 of Tab. IV. The expressions are parameterized by maximum-length parameter λ and/or maximum-gap parameter γ. We used the NYT dataset and ran FSM for different configurations of increasing difficulty w.r.t. output size. The results are shown in Fig. 4(a) using log-scale. For n-grams (first two groups), we observed that DESQ-DFS was up to two orders of magnitude faster than cSPADE. We only show the result for our own cSPADE implementation; the original C++ implementation was significantly slower. For example, for mining 10% of NYT, the original cSPADE implementation took more than 3 hours whereas our implementation took 400 seconds. Both cSPADE implementations were significantly slower than prefix-growth and DESQ-DFS, however, because cSPADE follows a candidate-generation-and-test approach and suffers from an excessive number of generated candidates. To keep our study manageable, we stopped cSAPDE after 12 hours. Compared to prefix-growth, DESQ-DFS had negligible overhead (less than 2.5%). For gap constraints (third and fourth group), DESQ-DFS was competitive and had an overhead of less than 10% over prefix-growth. This overhead is expected as pattern expressions for gap constrains have uncaptured wildcards (cf. T_2 in Tab. IV), which increases nondeterminism in the corresponding cFSTs and thus leads to more snapshots. For generalized n-grams (last group), where we additionally considered item hierarchies, the overhead was slightly more pronounced (up to 13%). Here the amount of backtracking performed by DESQ increased with the depth of hierarchy (cf. line 24 of Alg. 1 and discussion in Sec. IV-B).

We also investigated the overhead in terms of memory consumption. The results are shown in Fig. 4(b). For cSPADE, we report the maximum size of the inverted index and for prefix-growth and DESQ-DFS, we report the maximum size of the projected database. For n-grams and gap-constraints, DESQ-DFS had an overhead of up 18% and for generalized n-grams up to 23%. The overhead is unavoidable as for DESQ-DFS, we need to store cFST snapshots compared to only positional information as in prefix-growth and cSPADE. We may, however, improve memory consumption by swapping projected databases to disk [11]. We also ran the above experiments for the AMZN dataset (not shown here) and observed a similar behavior.

RE constraints. In this set of experiments, we evaluated the efficiency of DESQ for mining frequent subsequences (all or contiguous) that match a RE. Our pattern expressions allows us to express REs with their equivalent pattern expressions (cf. Tab I and expressions P_1–P_4 of Tab. IV). We compared DESQ's performance against state-of-the-art RE-constraint FSM methods SMA and prefix-growth. We used the PRT dataset and obtained suitable RE constraints from the PROSITE database [26]; the runtimes are shown in log-scale in Fig. 4(c). We observed that DESQ was up to 2.5x slower than SMA for P_1 and up to 1.3x slower than SMA on P_2. We do not give SMA results for P_3 and P_4 because the implementation produced incorrect results (acknowledged by the original authors). We did not investigate this further as the

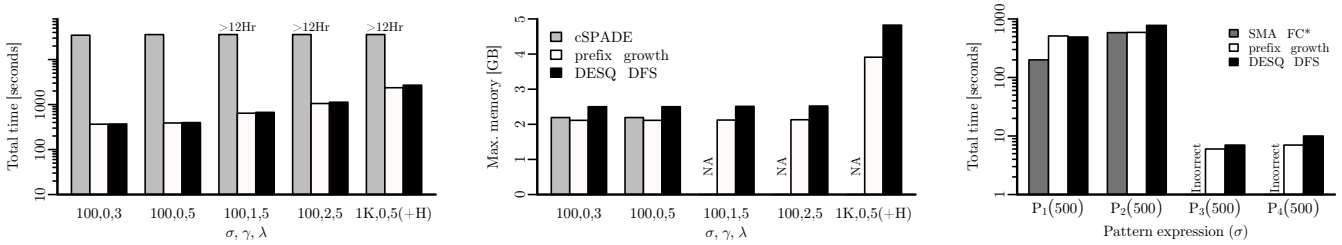

(a) NYT (length, gap, hierarchy)-constraints (b) NYT (length, gap, hierarchy)-constraints (c) RE-constraints

Fig. 4: Overhead of DESQ for common subsequence constraints.

SMA source is not available. DESQ was roughly on par with prefix-growth for P_1–P_4 (up to 1.3x) slower. The overhead of DESQ thus appears acceptable.

C. DESQ for Customized Subsequence Constraints

We evaluated the performance of DESQ for customized subsequence constraints that may arise in applications.

Constraints. We considered pattern expressions that express constraints in information extraction (IE), natural language processing (NLP), and customer behavior mining applications. These expressions are shown in Tab. IV. Expressions N_1–N_5 express constraints useful for IE and NLP applications and are inspired from [4], [5], [27], [28]; these expressions were used on the NYT dataset. Expressions A_1–A_4 expresses constraints useful for market-basket analysis and apply to AMZN.

DESQ algorithms. We evaluated the performance of Naïve, DESQ-COUNT and DESQ-DFS on pattern expressions N_1–N_5 and A_1–A_4. The results are shown in Fig. 5, which also gives the minimum support threshold σ used for each pattern expression (chosen empirically). The runtimes are given in log-scale. On the NYT dataset, for expressions N_1–N_3, DESQ-COUNT and DESQ-DFS had similar performance and finished in a few minutes. For N_4–N_5, however, runtimes were higher and DESQ-DFS was significantly faster than DESQ-COUNT (up to 14x). To gain insight into these results, we computed the average number μ of P-sequences (average of $|G_P^F(T)|$).[3] These numbers are shown above the bars for each pattern expression. We observed that for small values of μ, DESQ-COUNT and DESQ-DFS had similar performance, whereas for larger values of μ, DESQ-DFS was much more efficient. When μ is small, the simple counting method of DESQ-COUNT is expected to work well because few sequences are generated. The advanced pruning methods of DESQ-DFS are then not needed. When μ is large, however, DESQ-COUNT can enumerate many sequences that turn out to be infrequent, which is expensive. DESQ-DFS prunes many of these sequences early on and is thus more efficient. On the AMZN dataset (expressions A_1–A_4) DESQ-DFS consistently outperformed DESQ-COUNT (up to 22x). This behavior is explained by the observation that μ was large for all pattern expressions.

Based on these results, we conclude that DESQ-DFS consistently worked well in our experiments. Although DESQ-COUNT was slightly faster in some cases, it blew up on others. Thus we consider it generally safer to use DESQ-DFS in practice.

[3]We averaged over input sequences T for which $G_P^F(T) \neq \emptyset$.

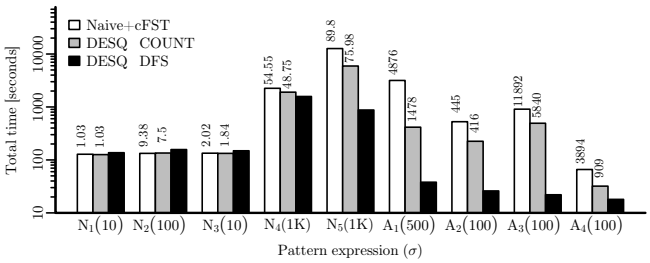

Fig. 5: Performance of DESQ mining algorithms. The numbers on top of the bars indicate the average number of P-subsequences per input sequence.

TABLE V: Runtimes of Naïve with cFST and openFST

	N_1	N_2	N_3	N_4	N_5	A_1	A_2	A_3	A_4
OpenFST	1h	1h	1hr	6h	>12h	>12h	>12h	>12h	50m
cFST	2m	2m	2m	30m	3.5h	1h	10m	15m	1m

D. FST Optimizations

We also compared cFST simulation with the state-of-the-art OpenFST library. We used the Naïve "generate-and-count" approach on pattern expressions N_1–N_5 and A_1–A_4. The results are shown in Tab. V. We observed that Naïve was orders of magnitude faster when used with cFST simulation than when used with OpenFST. This is because pattern expressions often translate to excessively large FSTs, which are inefficient to simulate (see Tab. II and discussion on cFSTs in Sec. IV-B). Moreover, OpenFST cannot directly handle hierarchies and, as discussed in Sec. IV-B, and many of our pattern expressions cannot be determinized. We conclude that cFST compression is effective.

VII. Related Work

Sequential pattern mining. The problem of mining frequent sequential patterns was introduced by Agarwal and Srikant [29]. Their Apriori algorithm follows a candidate-generation-and-test approach. The subsequent GSP algorithm [1] exploits the antimonotonicity property of sequential patterns to efficiently generate and prune candidate sequences. SPADE [15] by Zaki also generates and prunes candidates, but it operates on an inverted index structure representation of the database. Pei et al. proposed the PrefixSpan algorithm [11], which is based on a more efficient pattern-growth approach that recursively grows frequent prefixes using database projections. DESQ-DFS can be seen as a generalization of PrefixSpan to support arbitrary pattern expressions. SPAM [30],

which is similar to SPADE, uses an internal bitmap structure for database representation and employs a pattern-growth approach to mine frequent sequential patterns. A comprehensive discussion of these methods is given in [31].

Subsequence constraints.. There are many extensions to the basic sequential pattern mining framework for supporting subsequence constraints. GSP [1] and LASH [8], for example, allow gap constraints and incorporate item hierarchies. cSPADE [13] handles length, gap and item constraints. Wu et al. [32] consider subsequences with periodic wild card gaps, i.e., subsequences where consecutive items are separated by the same gap in the input. RE constraints have been studied by [9], [12], [17], [18]; these methods do not support capture groups. Some of the above constraints (e.g., gap constraints) target the input sequence, whereas others (e.g., length constraints, RE constraints) target subsequences. Pattern expressions unify both targets and allows us to express all of the above subsequence constraints (e.g., see Tab. I) as well as customized subsequence constraints that arise in FSM applications (e.g., see Tab. IV).

Pattern matching. Our work is also related to to pattern matching. There are many languages and systems for pattern matching over sequences; e.g., for information extraction [33], [34], computational linguistics [35], complex event processing [36], and sequence databases [37], [38]. Our pattern expressions are simpler than most pattern matching languages, yet expressive enough to specify many subsequence constraints that arise in applications. Nevertheless, pattern matching languages can conceivably be used to specify subsequence predicates and mine P-frequent sequences using Naïve, i.e., by first enumerating all matches and subsequently counting frequencies. Our experiments indicate that this approach is infeasible for many subsequence constraints. Instead, it is beneficial to integrate pattern matching and mining, e.g., along the lines of DESQ-COUNT and DESQ-DFS. An interesting direction for future work is to investigate to what extent such integration is possible for more powerful pattern matching languages.

VIII. Conclusions

In this paper, we introduced subsequence predicates as a general model for unifying and extending subsequence constraints for FSM. We proposed pattern expressions as a simple, intuitive way to express subsequence constraints, and suggested compressed finite state transducers as an underlying computation model. We also proposed the DESQ-COUNT and DESQ-DFS algorithms for efficient frequent sequence mining. Our experiments indicate that DESQ is an efficient, general-purpose FSM framework for common as well as customized subsequence constraints.

References

[1] R. Srikant and R. Agrawal, "Mining sequential patterns: Generalizations and performance improvements," in *EBDT*, 1996, pp. 3–17.

[2] J. Srivastava, R. Cooley, M. Deshpande, and P.-N. Tan, "Web usage mining: Discovery and applications of usage patterns from web data," *SIGKDD Explor. Newsl.*, vol. 1, no. 2, pp. 12–23, 2000.

[3] A. Lopez, "Statistical machine translation," *ACM Comput. Surv.*, vol. 40, no. 3, pp. 8:1–8:49, 2008.

[4] A. Fader, S. Soderland, and O. Etzioni, "Identifying relations for open information extraction," in *EMNLP*, 2011, pp. 1535–1545.

[5] N. Nakashole, G. Weikum, and F. Suchanek, "Patty: A taxonomy of relational patterns with semantic types," in *EMNLP-CoNLL*, 2012, pp. 1135–1145.

[6] A. Brazma, I. Jonassen, J. Vilo, and E. Ukkonen, "Pattern discovery in biosequences," *LNCS*, vol. 1433, pp. 257–270, 1998.

[7] K. Beedkar, K. Berberich, R. Gemulla, and I. Miliaraki, "Closing the gap: Sequence mining at scale," *ACM Trans. Database Syst.*, vol. 40, no. 2, pp. 8:1–8:44, 2015.

[8] K. Beedkar and R. Gemulla, "Lash: Large-scale sequence mining with hierarchies," in *SIGMOD*, 2015, pp. 491–503.

[9] M. N. Garofalakis, R. Rastogi, and K. Shim, "Spirit: Sequential pattern mining with regular expression constraints," in *VLDB*, 1999, pp. 223–234.

[10] I. Miliaraki, K. Berberich, R. Gemulla, and S. Zoupanos, "Mind the gap: Large-scale frequent sequence mining," in *SIGMOD*, 2013, pp. 797–808.

[11] J. Pei, J. Han, B. Mortazavi-Asl, H. Pinto, Q. Chen, U. Dayal, and M. Hsu, "PrefixSpan: Mining sequential patterns by prefix-projected growth," in *ICDE*, 2001, pp. 215–224.

[12] J. Pei, J. Han, and W. Wang, "Mining sequential patterns with constraints in large databases," in *CIKM*, 2002, pp. 18–25.

[13] M. J. Zaki, "Sequence mining in categorical domains: Incorporating constraints," in *CIKM*, 2000, pp. 422–429.

[14] K. Berberich and S. Bedathur, "Computing n-gram statistics in mapreduce," in *EDBT*, 2013, pp. 101–112.

[15] M. J. Zaki, "Spade: An efficient algorithm for mining frequent sequences," *Mach. Learn.*, vol. 42, no. 1-2, pp. 31–60, 2001.

[16] H. Mannila, H. Toivonen, and A. I. Verkamo, "Discovery of frequent episodes in event sequences," *Data Mining and Knowledge Discovery*, vol. 1, no. 3, pp. 259–289, 1997.

[17] H. Albert-Lorincz and J.-F. Boulicaut, "Mining frequent sequential patterns under regular expressions: a highly adaptative strategy for pushing constraints," in *SDM*, 2003, pp. 316–320.

[18] R. Trasarti, F. Bonchi, and B. Goethals, "Sequence mining automata: A new technique for mining frequent sequences under regular expressions," in *ICDM*, 2008, pp. 1061–1066.

[19] M. Mohri, "Finite-state transducers in language and speech processing," *Comput. Linguist.*, vol. 23, no. 2, pp. 269–311, Jun. 1997.

[20] K. Thompson, "Programming techniques: Regular expression search algorithm," *Commun. ACM*, vol. 11, no. 6, pp. 419–422, 1968.

[21] M. O. Rabin and D. Scott, "Finite automata and their decision problems," *IBM J. Res. Dev.*, vol. 3, no. 2, pp. 114–125, 1959.

[22] The New York Times corpus, http://catalog.ldc.upenn.edu/LDC2008T19.

[23] Web data: Amazon reviews, snap.stanford.edu/data/web-Amazon.html.

[24] SMA, http://www-kdd.isti.cnr.it/SMA/.

[25] OpenFST, http://www.openfst.org.

[26] PROSITE, http://prosite.expasy.org/.

[27] Google n-Grams, https://books.google.com/ngrams/.

[28] L. Del Corro, A. Abujabal, R. Gemulla, and G. Weikum, "Finet: Context-aware fine-grained named entity typing," in *EMNLP*, 2015, pp. 868–878.

[29] R. Agrawal and R. Srikant, "Mining sequential patterns," in *ICDE*, 1995, pp. 3–14.

[30] J. Ayres, J. Flannick, J. Gehrke, and T. Yiu, "Sequential pattern mining using a bitmap representation," in *SIGKDD*, 2002, pp. 429–435.

[31] N. R. Mabroukeh and C. I. Ezeife, "A taxonomy of sequential pattern mining algorithms," *ACM Comput. Surv.*, vol. 43, no. 1, pp. 3:1–3:41, 2010.

[32] Y. Wu, L. Wang, J. Ren, W. Ding, and X. Wu, "Mining sequential patterns with periodic wildcard gaps," *Applied intelligence*, vol. 41, no. 1, pp. 99–116, 2014.

[33] R. Krishnamurthy, Y. Li, S. Raghavan, F. Reiss, S. Vaithyanathan, and H. Zhu, "Systemt: A system for declarative information extraction," *SIGMOD Rec.*, vol. 37, no. 4, pp. 7–13, 2009.

[34] D. E. Appelt and B. Onyshkevych, "The common pattern specification language," in *TIPSTER*, 1998, pp. 23–30.

[35] O. Christ, "A modular and flexible architecture for an integrated corpus query system," *CoRR*, vol. abs/cmp-lg/9408005, 1994.

[36] N. Dindar, B. Güç, P. Lau, A. Ozal, M. Soner, and N. Tatbul, "Dejavu: Declarative pattern matching over live and archived streams of events," in *SIGMOD*, 2009, pp. 1023–1026.

[37] P. Seshadri, M. Livny, and R. Ramakrishnan, "The design and implementation of a sequence database system," in *VLDB*, 1996, pp. 99–110.

[38] E. Lo, B. Kao, W.-S. Ho, S. D. Lee, C. K. Chui, and D. W. Cheung, "Olap on sequence data," in *SIGMOD*, 2008, pp. 649–660.